

Summary and Analysis of A13

An Overview of the Technical Goals and Concepts in "*An Open Architecture for a Digital Library System and a Plan for its Development*" by Robert E. Kahn and Vinton G. Cerf (1988)

by Mark Stefik

Research Fellow
Palo Alto Research Center
Palo Alto, California 94304

Prepared for ContentGuard
May 30, 2007

Contents

Executive Summary.....	3
1 Introduction.....	5
2. Setting Goals for Digital Libraries.....	7
2.1 Setting Goals for Regulated Access.....	8
2.2 Setting Goals for Commercial Use.....	10
3. Technology Choices.....	14
3.1 Trust among Multiple Parties.....	16
3.2 Communication and Code Security.....	19
3.3 Representing Agreements.....	20
4. Concluding Remarks.....	22
Appendix A. Qualifications.....	23
Education and Research.....	23
Relationship to Robert Kahn and Vinton Cerf.....	24

Executive Summary

In 1988 Robert Kahn and Vinton Cerf wrote a draft paper (A13) proposing an architecture and plans for a digital library system. In broad terms, they proposed an infrastructure intended to link libraries together and to enable people to share information and documents.

In setting out their proposal, they made important design choices. A crucial choice at the onset was to base their design on the metaphor of a traditional library. In contrast, digital rights management systems such as those invented in the early 1990s and deployed today were intended to support digital commerce. These DRM designs were based on the metaphor of the electronic marketplace¹. The difference in choice of design metaphor led to different goals for what the systems were expected to do, leading in many crucial ways to essentially opposite design choices.

Traditional libraries in the United States are not commercial organizations. Traditional libraries are not concerned with regulating the particular uses to which information is put. Nor are they concerned with meeting the needs of commerce, such as pricing regimes, licensing arrangements, sales territories, or special deals for students or members of various groups. The most important legal basis for libraries is Copyright Law. In contrast, the legal basis for business agreements between parties is Contract Law.

This difference in orientation between A13 and later DRM systems is profound. In the main, A13 does not focus on the needs of commerce, where parties are often competitive and sometimes adversarial. For example, A13 does not consider any need for visible agreements reflecting the interests of parties in commercial transactions. It does not recognize a need for a threat analysis and arrangements for robust security in system architecture.

Table 1 summarizes some of the main differences in design choices between A13 and many later DRM approaches.

¹ For example, see the book *Internet Dreams: Archetypes, Myths, and Metaphors* by Mark Stefik (with a foreword by Vinton Cerf (one of the authors of A13)), published in 1996 by The MIT Press.

Design Issue	Design Choice in A13	Design Choice in DRM systems
Main design goal.	Promote sharing of information. A13 was intended to provide friendlier services akin to traditional libraries.	Promote a viable system to enable new business models for the commercial distribution of digital content.
Mechanism for overseeing rights.	Software-only agents ("Knowbots") that are attached to documents and travel from system to system as mobile code. Potentially, each different set of requirements is met by a different program.	Digital rights may be expressed in a declarative rights language for digital contracts. These contracts are intended for use both in user interfaces and by standard and certified trusted systems that enforce usage obligations. Communication among trusted systems is often protected by encryption. The trusted systems get the parameters of usage from the digital contracts. The same trusted systems potentially work for all documents. Foundations for trust in these systems may include physical (hardware) security, communication security (e.g. encryption), and behavioral security (certified programs).
Organizational basis of trust.	Content consumers must trust the producers of content, who commission the writing of different Knowbots for each use. Content producers must trust the consumers of content, who configure the operating environments on their machines for Knowbots. There is no consideration of the practical verifiability of system correctness or security.	Producers and consumers of content agree on terms and conditions and express them in digital contracts. They do not depend on each other's code. Rather, they rely on qualified (and disinterested) third parties to develop and certify trusted systems that enforce the digital contracts.
Purpose of computer oversight.	Knowbots are intended to represent the interests of content owners. Knowbots are largely concerned with accounting for the amount of usage for each document.	Digital contracts are designed to represent <i>all</i> essential elements in the agreements between all of the parties to the transactions. Digital contracts address the terms and conditions of use – including allowed operations, fees, timing, special licenses,

		distribution, and so on.
Accommodation of existing systems.	A13 recognized that the digital library would be distributed, heterarchical, networked, and display oriented. It must have an ability to interact with Digital Library Systems that do not adhere to its internal standards and procedures.	Most DRM systems invented in the 1990s and deployed today were a departure from existing networked file systems in their foundations of trust. In order to prevent the compromise of commerce or private content, they were designed specifically to only operate with other trusted systems that could meet their standards.

The rest of this paper discusses the technical goals and choices in A13 in more detail, and shows how the design choices took opposite paths from the DRM systems that were invented in the 1990s.

1 Introduction

In 1988 Robert Kahn and Vinton Cerf wrote a visionary white paper, “An Open Architecture for a Digital Library System and A Plan for its Development” (A13). It presents a draft research and development plan for a public information infrastructure that would enable digital library services. Kahn and Cerf were not newcomers to the creation of public infrastructure. They have received prestigious awards² for their earlier leadership and central roles in the creation and early development of the Arpanet – which famously became the Internet. Their recognized technical contributions to the Internet were for the protocols for packet-switching – the TCP/IP protocols. These transport protocols are called “low-level” because they govern how computers robustly send bits to each other in a digital network. With their backgrounds in computer science and electrical engineering, and networking specifically, Kahn and Cerf were well prepared to define these protocols. From their positions in the Information Processing Techniques Office of ARPA, they were deeply connected to the major centers of computer science research at the time, and well positioned to guide the creation of the net.

² Among other awards, Robert Kahn and Vinton Cerf were the winners of the Turing Award in 2004 for their pioneering work on networking. In 2005 they were awarded the Presidential Medal of Freedom for this work. They were inducted into the National Inventors Hall of Fame in 2006.

In A13, Kahn and Cerf considered a new challenge. They chose the metaphor of the “library” to set goals and expectations for a proposed higher-level digital information infrastructure.

The term “library” conjures a variety of different images. For some, a library is a dim and dusty place filled with out-of-date texts of limited historical interest. For others, it is a rich collection of archival quality information which may include video and audio tapes, disks, printed books, magazines, periodicals, reports and newspapers. As used in this report, a library is intended to be an extension of this latter concept to include material of current and possibly only transient interest. Seen from this new perspective, the digital library is a seamless blend of the conventional archive of current or historically important information and knowledge, along with ephemeral material such as drafts, notes, memoranda and files of ongoing activity. [A13, Summary, page 3]

The theme of Kahn and Cerf’s earlier, seminal contribution to the Internet was to “link computers together,” creating an infrastructure for sharing data. In A13, they develop the related theme to “link libraries together” in an infrastructure that would enable people to share information and documents. In their view, users would participate with their personal computers acting as personal digital libraries.

In its broadest sense, a DLS is made up of many Digital Libraries sharing common standards and methodologies. It involves many geographically distributed users and organizations, each of which has a digital library which contains information of both local and/or widespread interest. [A13, page 3]

Ultimately, the success and scope of A13 were limited by the library metaphor itself, which has proven unsuitable for fully characterizing commerce and information sharing in human organizations. This problem is that the normal activities in information exchange among competitive and sometimes adversarial human organizations are outside the scope of what normally happens in libraries. When A13 left behind the relatively simple world of computers that exchange bits, it entered the complex world of human organizations where activities involve ownership, competition, collaboration, and variations in law and agreement. This broader world of human commerce has many complexities and requirements that are not significant in libraries. Nor are these requirements comprehended or addressed in the goals and technologies described in A13.

- **A13 Goals.** In the U.S., the most influential libraries are public libraries and university libraries. The needs and experiences of these cooperative, non-

commercial institutions have proven to be somewhat misleading when used to set goals for a large-scale, commercially-oriented information infrastructure. For example, the library-oriented vision lacks a realistic consideration of information security and regulation of how information is used, which play central roles in the activities of online commercial systems.

- **A13 Technology.** For their leadership of the digital library project, Kahn and Cerf had to reach well beyond the EE/CS research that they knew best. For example, the technical projections on which elements of their approach depended (such as progress in natural language understanding by 2003) have proven to be overly optimistic at least in time-scale. Furthermore, some of the main technical ideas in their proposal have so far proven unworkable, and have been superseded by very different approaches. For example, A13 suggests transmitting mobile code (“Knowbots”) on behalf of a user to run at distant sites in order to search through their information files. The technology approaches that have come to dominate searching services in the World Wide Web are a study in contrast to A13’s proposals. For example, Web search services such as Google employ web-crawlers and indexing engines with massive resources³ for computation and storage⁴. Restated, the technical approaches that have proven most practical for massive search transmit content rather than transmitting programs (“Knowbots”).

The issues around goals and technologies in A13 are elaborated in the following sections.

2. Setting Goals for Digital Libraries

In the United States, most major libraries are either public libraries or are affiliated with higher educational institutions. Although there are also corporate research libraries and some collections in private ownership, most of these libraries are connected to academic libraries, especially for purposes of obtaining access to collections through inter-library loans. Kahn and Cerf also mention databases – for information such as scientific data, public records, law records and medical data. Such databases are mostly outside of the

³ Computation facilities known as “server farms” are employed by search companies to index the web. Server farms can cover several acres and have upwards of tens of thousands of computers.

⁴ Although web sites can have their own search services, their information is part of what is sometimes called the “dark web” if the information is not open for search by the main search engines.

library system and have not been freely accessible to the public. In their proposed “digital library project,” however, Kahn and Cerf sweep all of these sources and kinds of information into the familiar library pattern.

The mainstream activities of public and academic libraries establish key expectations. For example, libraries are not organized as commercial enterprises and their services are not for hire. Public libraries serve a public good, and academic libraries prioritize the particular needs of their academic communities. Libraries do not charge for information. Monetary transactions in libraries are mainly about record keeping and cost recovery for inter-library loans. Libraries do not advertise or suggest that people should buy certain sources of information. Librarians often provide facilities for making free copies of materials and are at least generally aware of copyright laws and the principle of fair use, whereby library clients may copy certain materials for their own scholarly purposes. Most libraries don’t house secret information such as private commercial data or data for which there are privacy concerns. Their goal is to make all of their information holdings available to all of their clients rather than making selected information differentially available to different people.

The mission of traditional libraries is to make information available to the public. In the main, libraries are not concerned with regulating who has access to information or what people can do with information once they have it. These library-centered assumptions about information and its use are reflected in the design goals and assumptions of the digital library proposal described in A13. Specifically, they affect how A13 sets expectations and goals about regulated uses, about commerce and business models, and about security.

2.1 Setting Goals for Regulated Access

In the world of competitive and adversarial commercial activities, organizations create trust boundaries that regulate access to information. For example, the accounting department in an enterprise keeps its files of information within a trust boundary where access is limited to people with particular authorizations. The ability to read or update accounting records is limited to those who have authorization to operate inside the trust boundary. A company often has multiple trust boundaries for different parts of its

operations. For example, human resource records are used by different people than engineering plans for future products. For another example, law firms representing multiple clients must keep the client information separate. Two competing companies do not generally have access to each others' customer, payroll, sales, or strategic planning information.

A diagram of the information infrastructure for a commercial environment would have elements supporting information use and elements supporting security. For example, it would include firewalls and other security systems to detect intrusions and to keep malicious software out. It would have cryptographic services that support codes to protect information from prying eyes and, as part of a communication subsystem, to robustly identify communicating parties and to protect the integrity of content. It would include authentication services to identify and authorize particular parties to access information.

Figure 1 reproduces a diagram from A13 showing the structure of the proposed Digital Library System. The figure shows that personal and organizational digital libraries are linked together on the Internet together with various databases. Boxes in the figure indicate services for registering documents, importing documents into the system, indexing and cataloging. There are also accounting and billing services – suitable for handling library fees. In the figure, multiple personal library systems co-exist with multiple organizational library systems without trust boundaries and few provisions for security⁵.

⁵ For example, page 23 of A13 describes a registration server as responsible for “registering new users, sources of information (databases) or other components newly added to the system.” In the further description about this, the main concern is about assigning unique identifiers as an aid to indexing and cataloging. No where are issues of verifying the authenticity of users or content discussed. Like an open library, the users are expected to behave in a kind of honor system. In contrast, more is at stake in an adversarial commercial environment. Digital systems supporting commerce and information access in such an environment would need robust and automatic means for checking the credentials of users and content before authorizing either access or changes to information.




Figure 1. Reproduction of Figure 2 from A13 – “Structure of the Digital Library System”

In understanding the scope and limitations of A13, what is interesting is what is missing. None of the normal commercial considerations involving security, authorization and trust are indicated in the diagram or elsewhere in the text of A13. There is no system “threat analysis” to characterize security requirements. Such provisions are outside the scope of libraries and they are outside the scope of the proposed digital library design⁶. Overall, A13 pays scant attention to commerce or security.

2.2 Setting Goals for Commercial Use

The early government-supported computer networks in the United States allowed no commerce. In the 1970s, if someone sent an email that could be construed as a “commercial message,” it would often set off a flurry of discussion about inappropriate use of the network. This “non-commercial” sensibility – coupled with the context of libraries – is part of the background context for A13.

Among the commercial businesses that use information, publishing stands out for its dependence on the regulation of uses of content in order to sustain its business. In publishing, authors and people in other creative roles create new content. Publishers and

⁶ On page 26 there is the following sentence: “To increase system integrity, the Accounting and Statistics Servers should be configured to accept data only from the appropriate sources and to raise alarms when data arrives from an unexpected source.” In this way the design focuses on accounting reliability, in contrast to security as it relates to regulated use of content.

distributors publish, distribute, and sell content. Consumers purchase and consume content for their information or entertainment purposes – such as playing music, watching a movie, or reading a book. People carry out different activities around information according to their roles as creators, owners, distributors and consumers. For example, consumers do not typically write in or modify books that they buy, make copies, and then distribute them in competition with the original authors or publishers.

Various arrangements are routinely employed to promote commerce in the sale and distribution of content. For example, content may be offered at a temporary discount to encourage sales. Tiered pricing is used to maximize profits by enabling different pricing in different markets. For example, price may depend on the time of sale, the location of the sale, or the affiliations of the purchaser. Special discounts may be offered for students, for members of a particular organization, for senior citizens, or for handicapped persons. For another example, businesses distinguish between the making of additional copies of content and the serial re-use of a single copy. Rental services circulate content without increasing the number of copies. Volume discounts are offered – reflecting differences in the costs of sales when large numbers of copies are sold at once. These sometimes socially-aware concepts arise in commercial publishing business practices. Successful use of these approaches in a digital publishing regime requires that information systems have adequate means for regulating the distribution and use of content.

In contrast, such an elaboration of requirements for success of commercial arrangements in a digital content regime is outside the scope of A13. A13 reflects the idea that a library is a place where clients get information. A13 has a model of owner's interests which is largely based on some accounting of levels of library usage *without* regulation of how information is used.

To their credit, Kahn and Cerf recognize that when information is digital, there are natural concerns about intellectual property protection. After all, compared to the substantial effort involved in reproducing (say) books as bound volumes, it is often very easy to copy (unprotected) digital information regardless of the amounts. Kahn and Cerf enumerate some of the problems and also acknowledge that solutions were not known at the time of their writing.

At present, the basis for intellectual property protection in the U.S. is Patent and Copyright law. The large scale aggregations of information found on CD-ROMs and the selective access to information found in on-line databases may require substantial re-thinking of the ways in which the creators and owners of such information are compensated for its use. There are many issues at stake in this area, not the least of which relate to the ease with which information can be replicated once in digital form and the rapidity with which large quantities of information can be processed (accessed, transferred, analyzed, integrated, etc.). Concepts of value and pricing and royalty for use of information could require considerable revision if the cost of such use is to remain within reason. One does not now pay an author a royalty each time a book is read. However, a royalty may be earned each time a song is played in public, though not in private. If a thousand books are combined on a single CD-ROM and the acquirer of the CD-ROM only intends to read one of them, what sort of royalty arrangement is appropriate to compensate the copyright owners? How would compensation be extended for cases in which electronic copies are provided to users? In fact, the concept of copying or duplicating a work may no longer be the essential factor in calculating royalties since far more complex actions may now be taken on digital information.

These questions are not trivial in nature nor have many workable solutions been proposed thus far. [A13, pages 11-12]

In this way, Kahn and Cerf encounter difficulties that are inherent in the library metaphor. Libraries do not themselves make copies of books (for example) in order to save on purchasing costs from publishers. They sometimes provide copiers so that clients can make their own limited copies of information for personal use. Such activities are governed by the provisions of U.S. copyright law. In the preceding quotation, Kahn and Cerf acknowledge that the kinds of activities likely to take place in a digital information infrastructure go beyond the routine activities in a library. When they suggest that "the concept of copying or duplicating a work may no longer be the essential factor in calculating royalties" they are in effect acknowledging that the economic agreements that form the basis of the publishing business break down when people make many (unauthorized and unrecorded) copies. In a digital network, provisions of Copyright Law apply, but enforcement can be intractable.

Kahn and Cerf recognized that solutions to this problem are needed. In terms of the legal basis for regulating use of content, they did not anticipate the possibility of using

Contract Law where Copyright Law falls short. Not recognizing the relevance of Contract Law, they also did not recognize the value of having a machine-understandable declarative language in which “contracts” can be expressed – so that their terms could be explicitly presented both to people and to computer systems. They did not anticipate that secure computer systems could play a practical role in the enforcement of the sort of declarative “digital contracts” that could robustly support the range of human agreements and activities typical in commerce.

In the 1990’s, other people developed concepts and technologies that addressed these problems⁷. These approaches led to today’s digital rights management (DRM) systems. In contrast to the library metaphor, DRM approaches recognize that there are different and distinguished uses or operations on content, such as making copies, modifying content, printing it, distributing it, selling it, and so on. DRM systems provide means for regulating these different uses in order to sustain an economy in which content is produced and consumed. Conditions are associated with the operations. These conditions must be satisfied before the operations can be performed. For example, a person might need to pay a fee, be a certain age, live in a certain jurisdiction, or belong to a particular group, and so on.

It is not surprising that these distinguished operations and conditions are outside the scope of the library metaphor. Libraries do not create or sell content. They are concerned mainly with making information available to the public. Once a client “gets” information from a library – such as checking out a book—the library is not involved with what the client does with it. Libraries generally do not regulate how information is used or who has access⁸.

⁷ See for example, Stefik, Mark. “Letting Loose the Light: Igniting Commerce in Electronic Publication.” In *Internet Dreams*, The MIT Press, Cambridge, Ma., 1996, pages 219-253. The foreword of this book was written by Vint Cerf, one of the authors of A13. The book proposes four metaphors for understanding the origins and directions of the Internet: digital libraries, electronic mail, electronic markets, and digital worlds. In contrasting digital libraries to electronic markets, it highlights the different capabilities that are suggested by the different metaphors. Specifically, this book elaborates how the metaphor of the electronic market usefully covers activities that are not associated with libraries.

⁸ The more recent controversies about regulating access by children to certain materials on the Internet, and the reporting of which people have had access to certain information, has been controversial in part because of the conflict with the over all mission of libraries to provide the public with access to information with a minimum of barriers.

Discussion of these issues in A13 adheres to the expectations set by libraries. When Kahn and Cerf imagine what people will do with the information in digital libraries, they say people will “register, store, catalog, search, retrieve, and manipulate digital information in the library” (A13, page 8). These are the same kinds of activities that people could already do in traditional libraries. Notably absent from this list are commercial activities like selling or distributing information, or re-publishing information in revised forms.

In summary, A13 does not investigate the commercial situations where different people have different roles and sanctioned activities around information. A13 characterizes operations on information in library terms – relating to catalogs and retrieval and so on. It does not analyze requirements of a publishing business or requirements among competitive enterprises around regulated use of information.

3. Technology Choices

A13 proposes a high-level draft architecture for a digital library infrastructure. Analogous to the network architecture of the Internet, Kahn and Cerf proposed a network of computers where the nodes of the network are digital libraries. The architecture also included specialized servers for indexing, cataloging, registration, and other specialized functions.

Figure 2 reproduces a diagram from A13 illustrating its high-level system elements for a personal library system (PLS)—a node in the network of libraries. The diagram shows the structure of the PLS in terms of “layers” with an operating system and its device drivers in the bottom layer and application elements in the top layer. The second layer from the bottom contains file services, presentation (display) services, and network transport services. The top layer of the diagram divides elements into ones that serve the user (user interface), ones that access library content, and administrative functions.

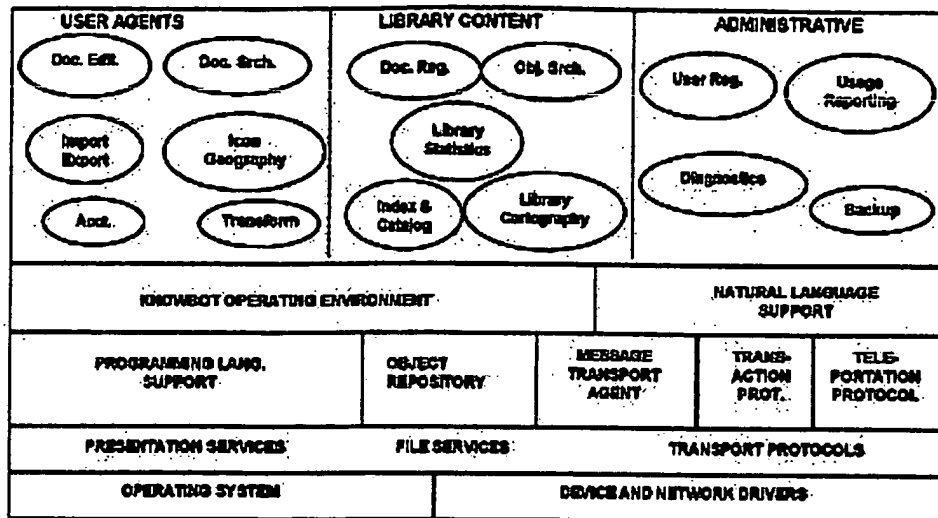


Figure 2. Reproduction of Figure 3 from A13 – “Personal Library System Structure”

The fourth layer of the architecture is of particular interest and contains the most novel aspects of the design. It contains a Knowbot operating environment (KNOE) and natural language support. Knowbots are central in Kahn and Cerf's design for many of the capabilities of the digital library infrastructure. A13 describes Knowbots as follows:

A Knowbot is an active program capable of operating in its native software environment. Knowbots are present in each of the various components of a Digital Library System. They can be cloned, replicated, created, destroyed, can be resident at a given host system or can move from one host machine to another. Knowbots communicate with each other by means of messages.

Knowbots act as the primary medium of communication and interaction between various major components of the Digital Library System. They may even transport other Knowbots. Generally, a Knowbot may be viewed as a user Knowbot or as a system Knowbot depending on whether it directly serves an individual user or not. [A13, page 34]

Knowbots are more than an incidental part of the proposal for a digital library infrastructure. Unlike the other system architectural concepts, A13 devotes a full section for describing Knowbots (Section 3: Knowbots and Their Application). As Kahn and Cerf noted:

The selection of a methodology for building Knowbots and even the determination whether an object-oriented language is essential are two

of the highest priority research questions for the Digital Library Project to resolve. [A13, page 31].

Many of the issues with the architecture in A13 arise from fundamental difficulties with Knowbots, especially security issues⁹. The use of Knowbots as mobile code in the system architecture was an aggressive design choice. Although concepts of object-oriented programming had been developed in the Computer Science community for over a decade¹⁰, the research focused on systems for single computers¹¹. Object approaches for distributed and mobile code applications were less explored. Nor had object-oriented applications been deeply explored involving mobile objects or their security considerations across multiple organizations.

The following sections examine Knowbot-related problems with the design in A13 involving trust among multiple parties, security vulnerabilities, and practical problems in using opaque ("black box") mobile code to represent agreements between producers and consumers of content.

3.1 Trust among Multiple Parties

Knowbots are employed in the proposed digital library architecture to carry out many functions. As described in A13, Knowbots can be asked to retrieve or file documents. They also watch over information objects on behalf of owners.

One set of system Knowbots specifically attend to locally available library information. They take requests from user Knowbots and actually retrieve the documents from storage (or conversely store them away). Another set of system Knowbots attend to background and administrative tasks such as diagnostics, backup and accounting. [A13, page 34]

⁹ Since the focus of A13 is on libraries rather than commercial organizations, it does not focus on security issues. Nor does A13 explore security requirements for Knowbots.

¹⁰ On page 30, A13 cites Smalltalk, Common Lisp, Common LOOPS and C++ as examples of existing object languages. Smalltalk and Common LOOPS were both largely developed at PARC where I work. I was one of the contributors to the LOOPS and Common LOOPS specifications at the time and was an active researcher in the object-oriented programming community.

¹¹ For example, see Stefik, M. Bobrow, D.G. Object-oriented programming: Themes and Variations. *AI Magazine* 6:4, pp. 40-62, Winter 1986. (Reprinted in Peterson, G.E. (ed), *Object-Oriented Computing*, Volume 1: Concepts, IEEE Computer Society Press, pp. 182-204, 1987. Also reprinted in Richer, M.H. (ed.) *AI Tools and Techniques*, pp. 3-45, Ablex Publishing Corporation, Norwood, New Jersey.)

In carrying out these functions, Knowbots would have to travel through the network of libraries. When we consider this architecture in the context of commercial activities, a problem of trust arises. Suppose that Company A has a computer system. A document and its Knowbot arrives from Company B. How does Company A know that the Knowbot can be trusted?

A13 only partially addresses the issue of trust. In the context of courier Knowbots, it says the following:

A class of trusted Knowbots called *couriers* have the special responsibility to look after selected objects on behalf of their authors or other owners of rights in the objects. [A13, page 34]

In the example, the courier is expected to look after the interests of Company B. The problem is that the Knowbot consists of mobile code that is being enabled to run on the computers of Company A. What assurances does Company A have about what the Knowbot will do? Will it correctly enforce the agreements that Company A has with Company B about use of the document? When the Knowbot reports back to Company B, will it also send back any additional information that compromises the business interests of Company A? Will it confine its activities to the document from Company B, or will it read or modify any other sensitive documents in the library of Company A? Even if the intentions of Company B are completely legitimate, what if there is a bug in the mobile code that inadvertently leads to commercial losses for Company A? After all, the coders of the Knowbot have presumably not been able to test the Knowbot on the computers of Company A. Allegorically, these scenarios are akin to “turning a fox loose in the hen house.”

Such concerns about mobile code are not unrealistic. The most familiar examples of mobile code today are computer viruses. Analogous to Knowbots, viruses travel from system to system “carrying out the wishes of their creators” – which are generally disruptive. Most businesses invest heavily in virus detection and firewalls in order to prevent unwanted infections that can compromise their computers and disrupt their

businesses. Even today nearly two decades since A13 was written, mobile code has had very limited use beyond support for user interfaces in web pages¹².

Nor are all of the risks associated with Company A. Suppose that the KNOE (Knowbot Operating Environment) in the computers of Company A has been altered. It could potentially separate a Knowbot from its document, and provide its own altered-Knowbot which acts differently. It could under-report the usage statistics in order to reduce fees. It could alter the Knowbot from Company B and send it on to Company C in a way that intercepts a "funding stream" to the coffers of company A rather than Company B. Allegorically, these scenarios are akin to "sending a lamb to a den of wolves."

A fundamental problem in A13's conception of Knowbots is that they are designed to serve the interests of *one* party. However, commercial transactions *inherently* involve multiple parties with differing interests, such as a producer and a consumer. A13 does not discuss any requirement for robust, verifiable and accountable trust across multiple parties. Furthermore, since programs and programming environments are so complex, the concepts of Knowbots and KNOEs does not appear to be a workable approach for achieving such security¹³.

DRM systems like those invented in the 1990s address multi-party security issues in a way more appropriate for commercial use: DRM systems like those invented in the 1990s use host computers configured as trusted systems, rather than as untrusted hosts. Each trusted system is designed and verified by a third party to guarantee to two contracting parties that the trusted system can be relied upon to act as an impartial agent to enforce the terms of any digital contract that has been agreed to by those parties. In this way, DRM approaches address multi-party security issues in a way more appropriate for commercial use.

In summary, the Knowbot architecture of A13 does not recognize that trust in a commercial setting requires a basis of trust across multiple parties. The Knowbot-based

¹² Even the use of mobile code in the limited context of "JavaScript" in web pages introduces security issues.

¹³ Even if the "source code" of Knowbots and KNOEs were made available to both parties, since there would be so much variation in Knowbots and environments it would not be practical to prove their correctness. It would also be difficult to prove (for example) that Knowbots were not transported to different environments than the ones that were tested.

architecture contains no provisions for this. As suggested by scenarios above, it appears that the Knowbot approach is not a sound basis for building such systems.

3.2 Communication and Code Security

Libraries generally do not often have active adversaries and their every day operations do not give much attention on security in communication. For example, when a library loans a client a book, there is typically no concern about whether the book has been altered to contain false information. In contrast, commercial organizations have private and critical communication. It is normal security practice to use encryption, redundant check sums, and other techniques in commercial communication systems to assure secrecy and integrity. For example, regular web users are familiar with providing passwords and with the “encrypted page” messages that they get in their web browsers when critical information is requested. Such concepts, however, are outside the activities of traditional libraries. Furthermore, they are not considered or even mentioned in A13.

The lack of attention to cryptography, certification and other means of securing communication has widespread implications for the system design and its applications. For example, consider again the transmission of mobile Knowbots. Knowbots are proposed for many functions in the digital libraries – used for retrieval and search, to billing, and even as couriers. When a Knowbot is being transmitted from one system to another, it is represented by bits in packets over the communications network. Today, several years of experience with computer viruses has made designers of computer systems aware of a myriad of possible “attacks” on communications systems¹⁴. By various means, communications can be intercepted, faked, copied, altered, and blocked. Lacking such understanding, A13 does not recognize that its Knowbots are vulnerable to being intercepted, faked, copied, altered, or blocked.

To restate the design issue, A13 proposes Knowbots as the security mechanism to act on behalf of owners. However, Knowbots are software-only entities. They are expected to travel over networks where they are vulnerable to modification, and to run on many

¹⁴ For example, see Chapter 3, “The Digital Wallet and the Copyright Box: The Coming Arms Race in Trusted Systems” in Stefik, Mark. *The Internet Edge: Social, Technical, and Legal Challenges for a Networked World*. Cambridge, Ma. The MIT Press, 1999. Robert Kahn and Vinton Cerf, the authors of A13, both provided endorsements of the book to the publisher. These appear on the back cover of this book.

different computer systems where they potentially could be modified. This leaves the foundations of security in A13 – the infrastructure that is supposed to look after the interests of content owners – profoundly vulnerable and insecure. In contrast, modern security designs have design features that address security in communication (communication integrity), authentication (behavioral integrity), and hardware (physical integrity).

A13 specifically proposes to leverage the Internet but to integrate the digital library infrastructure with existing technologies.

Before describing specific features of the Digital Library System, it will be helpful to review some of the fundamental assumptions which strongly affect its design. Perhaps the most dominant of these assumptions are that the system is distributed, heterarchical, hierarchical, networked and strongly display-oriented. In addition, it must have an ability to interact with other autonomous Digital Library Systems that do not adhere to its internal standards and procedures.
[A13, page 19]

From a computer security point-of-view, a chain is as strong as its weakest link. Every subsystem that does not meet security standards creates a vulnerability. This is the opposite of the trusted system approach of DRM systems of the 1990s, which require that every system in a transaction be certified as trustworthy.

In summary, the approach taught in A13 is vulnerable to communications attacks which would need to be addressed in a context of commercial use.

3.3 Representing Agreements

When people engage in financial transactions, they often formalize their agreements with contracts. For example, contracts and warranties are common when people rent an apartment, buy a home, take out a loan, or buy an expensive appliance. There are many variations in contracts – in terms of what provisions are included. At the same time, there are often a few key issues, such as the loan amount, fees, and the interest rate and term of a loan. Contracts explain the terms and conditions of the agreements, and lay out the rights of the parties. The substance of an agreement is not hidden. Rather it is written openly in the contracts. Contracts are intended to express the terms and conditions so that any of the parties to a transaction can examine them and understand them.

Contracts do not play a highly visible role in traditional libraries. For example, there are very few variations in the terms and conditions of book loans. There may be some books that cannot be checked out, and there may be some books that have to be returned more quickly than others. However, contracts are not a major focus for a traditional library. When Kahn and Cerf chose the library metaphor for their proposed digital library infrastructure in A13, they did not mention contracts as a design element.

A13 proposes Knowbots to take the responsibility to look after the interests of users and others in the digital library. Absent any other means, Knowbots are A13's only vehicle for representing and enforcing agreements. When there are multiple, different agreements between parties, multiple, different Knowbots would be required.

The problem with this approach is that Knowbots are mobile computer code and computer code is notoriously non-transparent and hard to understand for most people. Most of the code in a program is generally about its internal bookkeeping and managing relations with other objects in its operating environment. To understand what code does, a programmer must understand not only the computer language, but also the operating environment in which the code is operating.

Using Knowbots to represent a myriad of business agreements is not practical. The important and salient elements of human agreements would be buried and essentially obfuscated by putting them into a program. The opaqueness would make it very difficult to understand what agreement is represented by a Knowbot, or even to tell the difference between a genuine Knowbot and a rogue Knowbot.

The DRM systems that were invented in the early 1990s addressed the problem of representing agreements through the use of a declarative digital rights language¹⁵. These languages were designed to express the kinds of operations, terms and conditions that are salient for practical contracts about the use of digital content. For example, there were specific operations for copying, loaning, printing and other common things. Terms and conditions could express a range of requirements for payments, times of use, and so on.

¹⁵ For example, see "Letting Loose the Light" in Stefik's *Internet Dreams* book for an example of a digital rights language. See "The Bit and the Pendulum" in Stefik's *The Internet Edge* book for a discussion of digital contracts.

Digital contracts could be expressed in a grammar. From there they could be presented to people in clear user interfaces. They could also be interpreted and enforced by computers.

In summary, the Knowbot approach is not practical for supporting the negotiation or understanding of agreements. The subsequently-developed rights language approach side-steps the main difficulties of the Knowbots approach by making it unnecessary for people to try to discern the meaning of an agreement from the programming code of a Knowbot.

4. Concluding Remarks

A13 was a visionary proposal for a digital library system. The design was guided by the metaphor of a traditional library as a system to enable people to share information. Like traditional libraries – A13's focus is on exchange of information. True to the purpose of libraries, A13 is not much concerned with an infrastructure for commerce in digital content, which would require much more attention to mechanisms for commerce and to security requirements for transactions among potentially competitive parties. A13 made very different (and often opposite) design choices from many of the DRM systems that were invented in the 1990's and the systems that are deployed today.

Appendix A. Qualifications

Education and Research

My university education was at Stanford University, both undergraduate and graduate. I received my Bachelors degree in Mathematics in 1970 and my doctorate in Computer Science in 1980. I am a Fellow in the American Association for the Advancement of Science (AAAS) and also in the American Association for Artificial Intelligence (AAAI).

I work at the Palo Alto Research Center (PARC), where I am a research fellow. Since I started at PARC in 1980 I have taken several tours of duty in research management, leading three technical areas and one of PARC's laboratories for several years. I occasionally teach courses and give lectures at Stanford University and U.C. Berkeley. I have been an external thesis advisor and dissertation committee member for Ph.D. students at Stanford, U.C. Berkeley, and the University of Maryland.

I have published five technical books including *The Internet Edge: Social, Technical, and Legal Challenges for a Networked World* (MIT Press, 1999), *Internet Dreams: Archetypes, Myths, and Metaphors* (MIT Press, 1996), and *Introduction to Knowledge Systems* (Morgan Kaufmann Press, 1995). I have published over forty technical papers.

Some of my technical work is mentioned in the *Open Architecture for a Digital Library System* paper (A13), which cites Common LOOPS among current object languages. I was one of the main creators of the Loops system and a contributor to the Common LOOPS specification. PARC was a center for much of the research that inspired Kahn and Cerf during this period – including the technologies for distributed computing, the SmallTalk language, and the NoteCards hypermedia system which are mentioned in the paper.

As a computer scientist, I am somewhat of a generalist and have switched my area of focus every few years. A unifying goal in my work has been to enhance the creation and sharing of knowledge. My dissertation work on an expert system for experiment planning included a frame-based knowledge representation system. A version of this was later commercialized by Intellicorp. My research on collaboration in electronic meeting rooms ("Colab") included creating an infrastructure for distributed objects. The Colab research

led to collaboration with Bob Kahn and others in the creating of the "National Collaboratory" projects in the U.S.

My current research on "sensemaking systems" is about technology to help people facing information overload to master and understand large amounts of information in carrying out their work. Our sensemaking projects at PARC are multi-disciplinary, involving computer scientists and cognitive psychologists as well as specialists in natural language technology, user interfaces, and distributed systems. The current directions in this involve what we call "augmented social cognition," which is the technology that aggregates and combines the sensemaking contributions of large groups of people.

Relationship to Robert Kahn and Vinton Cerf

Bob Kahn was a fairly frequent visitor to the Heuristic Programming Project at Stanford University in the mid 1970s when I was a graduate student there. I got to know him since he worked closely with Edward Feigenbaum, who was one of my faculty advisors. In the early 1980's after I had graduated, I was a participant together with about a dozen others in some weekend workshops led by Bob Kahn at Stanford when the early ideas for a "digital library project" were being developed. Some of the other participants at the workshop were professors from Stanford (Edward Feigenbaum, Joshua Lederberg, John McCarthy) and MIT (Marvin Minsky). In this way I was able to contribute in a small way to the early stages of the Digital Library Project, even before the Open Architecture paper was written.

Bob and I have worked together on various projects over the years. I typically see him two or three times a year. When I first developed the concepts of digital rights management in the early 1990s, Bob signed a non-disclosure agreement with the Palo Alto Research Center so that we could discuss the ideas in some depth and also plan participation in some coordinated activities, such as the "digital object identifier" project.

Vinton Cerf was a professor at Stanford in the Computer Science Department when I started there as a graduate student. Initially, my research plan was to do a dissertation in the Systems area and he was my graduate advisor. Within a year or so, however, he decided to leave Stanford to work on developing the Internet (then ARPANET) at DARPA. I switched research areas in Computer Science to artificial intelligence,

focusing on expert systems. We have kept in occasional contact over the years. When I published the book *Internet Dreams* with MIT Press in 1995, he graciously wrote the foreword to the book. This book contained my first publication (beyond patents) of the ideas for digital rights management in the chapter titled "Letting Loose the Light: Igniting Commerce in Electronic Publication." The book also includes an excerpt from Kahn and Cerf's paper *The World of Knowbots*.

In 1999 I wrote a second Internet book – *The Internet Edge* – that provided a deeper analysis of the trends in networks, and discussed social, technical, and legal challenges. Both Robert Kahn and Vinton Cerf provided endorsements of the book to the publisher that appeared on the jacket of the book.

Mark Stefik 30 May 2007